DOCUMENT RESUME

ED 155 643 CS 004 161

AUTHOR LaBerge, David
TITLE The Perception of Units in Beginning Reading.
INSTITUTION Pittsburgh Univ., Pa. Learning Research and Development Center.
SPONS AGENCY National Inst. of Education (DHEW), Washington, D.C.
PUB DATE Jun 76
CONTRACT 400-75-0049
NOTE 55p.; Paper presented at the Conference on Theory and Practice of Beginning Reading Instruction, University of Pittsburgh, Learning Research and Development Center, June 1976, For related documents see, CS 004 132-133, CS 004 135, CS 004 137-173, ED 125 315 and ED 145 399

EDRS PRICE MF-$0.83 HC-$3.50 Plus Postage.
DESCRIPTORS *Beginning Reading; Conference Reports; Context Clues; Decoding (Reading); *Learning Theories; *Pattern Recognition; *Perception; Primary Education; Reading Instruction; Reading Research; Reading Skills; Structural Analysis; *Visual Perception; *Word Recognition; Word Study Skills

ABSTRACT

Perceptual units are involved in the reading process in a particular way. Reading starts with patterns of visual perception visual stimuli in combination with either previously learned information or a strategy for selectively attending to part of the stimuli. Whether the visual pattern selected for deeper information processing is a letter, a spelling pattern, or a word group depends not only on the stimulus display (perceptual unit) but also on the dominant contextual mechanism. Research indicates that these mechanisms are driven by "top-down" hierarchical operations on patterns, phonological events, and semantic/syntactic events; the more sophisticated learning strategies are used for more sophisticated tasks to keep processing time at a minimum. Six models of processing strategies have been used to explain the operations on perceptual units, and their interchangeability suggests that reading strategies specially for word recognition and decoding might be learned in sequence and then used for specific reading tasks. The way individuals learn to recognize the first perceptual units at the new processing levels (new contextual mechanisms) remains a difficult issue that has yet to be resolved. (Discussion following presentation of the paper is included.) (RL)

**********
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
**********

U S DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

ED155643

# The Perception of Units in Beginning Reading

David LaBerge

University of Minnesota

This paper was presented at the conference on Theory and Practice of Beginning Reading Instruction held at the University of Pittsburgh, Learning Research and Development Center, June 1976.

Do not quote without permission. Comments would be appreciated.

Conferences supported by a grant to the Learning Research and Development Center from the National Institute of Education (NIE), United States Department of Health, Education, and Welfare, as part of NIE's Compensatory Education Study. The opinions expressed do not necessarily reflect the position or policy of NIE, and no official endorsement should be inferred. NIE Contract #400-75-0049

CS004 161

## The Perception of Units in Beginning Reading

The main concern of this chapter is the particular way that perceptual units are involved in the reading process. My remarks are not aimed at describing a complete theory of reading, nor even a theory of beginning reading. Rather I would like to treat a restricted aspect of reading which asks questions about units, their functions, their sizes, and their development in the experiences of the person. As it turns out, my remarks may provide more questions then answers; but I would be content if, in the course of this paper, one or two basic issues in this area were clarified so that a set of productive hypotheses could be generated for further research.

Most of the examples which will be treated here are related less to comprehension than to the perceptual aspect of reading, for example, word recognition and decoding, which are important issues in beginning reading. However it would be a bonus if some of the principles discussed at the perceptual end of the reading process would generalize to comprehension processes deeper in the cognitive system.

The plan of this paper falls into three parts. First I will discuss what I consider to be the main function of units for perceptual processing in general and for reading in particular. This will involve a description of a class of models which interrelate these units. The second part of the chapter will be concerned with the question of how we manage to select the particular set of units appropriate to the task at hand. Finally I will discuss in a general way some important factors which may influence the learning of new units.

### The Function of Units

It is rare that we hear anyone talk about reading without mentioning units. Certainly our conventions of measurement require that we agree upon

the definition of a unit. In measuring reading speed, we usually count words. In measuring immediate memory span, we may count letters; and when we measure a line of poetry we may count syllables. However, many psychologists consider words, syllables, and letters as more than simply objective units of measurement. They also regard them as conceptual representations of perceptual events or structures (Melton & Martin, 1972). They assume that the subject has a network or code representation for each of the many familiar patterns he has experienced, ranging from simple line-features to more complex patterns such as letters, words and faces. These codes may be considered as part of long-term memory; to be more precise, as long-term perceptual memory, and it is these structures which are first encountered as information comes off the sensory surfaces of the eye and ear. For example, when we are shown a new word, e.g., the name of a new flower, which I will call a sidularium, not everything we see as we look at it is new. The letters and probably the letter clusters are familiar enough so that we can spell the word and pronounce it. We may even try to relate the word to other things, or rehearse it so that we can recall it later. But before we do any of these things, the word stimulus falling on our retina has already made contact with the memory of certain letters and letter-clusters, and it seems as if the recognition of these familiar patterns simply happens to us. Therefore, this point of view introduces perceptual units very early in the processing system, and the outputs from these perceptual units are used to process higher order operations, such as naming, rehearsal, and comprehension.

Let us look at a schematic representation of perceptual units related to reading based on the assumptions set forth in the LaBerge and Samuels model (1974).

---

Insert Figure 1 about here

---

In Figure 1 there are three major systems which contain an array of codes or units: the visual perceptual system, the phonological system, and the semantic-syntactic system. I will use the terms unit and code interchangeably in this paper. Graphemic information from the printed page enters from the top, and makes contact with (i.e., activates) the codes of the visual system, represented by dots. Solid dots refer to codes which are well-formed in long-term memory through many experiences with a particular stimulus pattern such as a common letter or word. Open dots refer to codes which are unfamiliar, and therefore have been only temporarily formed, for example, a Greek letter or a new flower name.

Although the visual codes are spatially arranged in a hierarchy in Figure 1, there is no specific indication of how they are interrelated here, hence the omission of any lines connecting codes with each other. For it turns out that there are several alternative models (LaBerge, 1976) by which the codes may be linked together. These models we will discuss in detail later in connection with learning new codes. But regardless of which of these models is assumed, outputs from many of these visual codes or units become associated with units in other systems. For example, we learn to name letters and words; and may give meanings to words not only from their phonological units but also directly from their visual units.

The name and sound units are represented by codes in the phonological system, and here we may also assume a hierarchical arrangement of codes as shown in Figure 1. For example, when a child sounds out a new word, "clam", he may visually code "cl" and "am", and by association produce the sounds /kl/

and /æm/ which he then blends into a higher order sound /klæm/. With practice he may learn to associate the whole word "clam" directly to the sound /klæm/.

Other associative links in Figure 1 are assumed to connect phonological codes to meaning codes (which a child typically has learned before he begins to read), and to connect visual codes directly to meaning codes. These lines of association are assumed to be activated by outputs from unitary codes. And herein lies one of the advantages of perceptual unitizing. If no unitization occurred before intersystemic associations were attempted, then a good deal of confusion would accompany an association. For example, assume that "cl" and "am" were separately associated directly to the word sound /klæm/. Then the child would say /klæm/ to any word which began with "cl" or ended with "am". Obviously, to avoid confusions such as this one, either the child must unitize the visual clusters into a visual word before learning the association, or else always rely on associating spelling patterns with the appropriate syllables. But this approach will run into the same problem when patterns such as "th" and "ough" are encountered. These letters must be unitized into spelling patterns before the appropriate sounds can be associated to them without confusion.

A parallel line of reasoning is assumed for the case in which sounds are associated to meanings and visual units are associated with meanings. It is difficult to imagine how intersystemic associations can effectively occur unless there are unitizations of the perceptual codes, i.e., the child must unitize the pattern to sound it or get its meaning. Therefore, units can be regarded as transformers of information; they take information from the sensory surface, directly or indirectly through other subordinate units, and put it into a unitary form to transmit to other major cognitive systems for further processing. This transformation is not merely a matter of convenience, but

rather somewhat of a necessity. Evidentally some associative transform must take place between the information falling on the sensory surface and the phonological and meaning systems, because the acoustical codes and the meaning codes assigned to this information are almost always arbitrarily related to the form of the physical object. A unit seems to be an appropriate form which perceptual information assumes in order that arbitrary associative relationships can be made in an efficient and relatively unambiguous manner.

Units also act as filters on incoming sensory information (Estes, 1974). If perceptual units were omitted from the design of the human reader, then the myriad details of information at the sensory surfaces would have to be processed at the deeper cognitive levels. This would reduce the efficiency of processing, especially in view of the limited capacity characteristics of processing to be discussed in more detail in the following paragraph.

The size of the unit may vary, depending upon the task demands. For ordinary reading purposes, the larger the perceptual unit, the better (Gibson & Levin, 1975). Some stages of processing have limits on the number of units of information that can be processed simultaneously. One of these systems is immediate memory, which is said to have an upper limit of $7 \pm 2$ units (Miller, 1955) or $5 \pm 1$ units (Mandler, 1967). Presumably, we use this kind of memory to keep in mind the first part of a sentence while we are perceiving the last part of a sentence so we can put the whole sentence together for comprehension. Another system which has a sharply limited capacity is attention, which may have a limit of only one unit upon which it can operate at any given moment, (Broadbent, 1958; Treisman, 1964; Moray, 1969) although with rapid shifting, attention may appear to maintain several units at a state of heightened activity at one time. Given that immediate memory and attention capacities are

indeed so small, the skilled reader evidently processes information in text as fast as he does because he packs more information into each unit.

As it turns out, the information—carrying capacity of perceptual units (and meaning units as well) can be very large, but it costs something to the reader, namely a considerable amount of learning, to compress information into a unit quickly and reliably. Perceptual learning is considered to proceed at a relatively slow rate as compared to associative learning which may often occur in an all-or-none fashion (Estes, 1970). Nevertheless, learning to perceive information in terms of large units is critical for fluent reading and for higher cognitive functions such as creative thinking. Therefore, it is not surprising that reading skill requires years to become fluent.

One of the benefits of unitizing perceptual patterns is that the processing usually becomes more automatic with each exposure. This means that less attention is needed to process a given unit, so that attention can be devoted to other aspects of the task at hand such as comprehension and coding for recall (LaBerge & Samuels, 1974). For example, recognizing a new word at first takes attentional effort, but later, as it becomes more familiar, seems to be recognized with no effort; in fact, we cannot easily stop ourselves from recognizing it. The relationship between familiarity and automaticity has been studied in experiments by Shiffrin and Gardner (1972) and by LaBerge (1973). Currently we are exploring other ways to measure automaticity and relate it to measures of unitizing during perceptual learning.

At this time, we feel relatively secure in our assumption that the familiar units or codes represented in Figure 1 are processed automatically from information falling on the sensory surface. Therefore, when a familiar word is shown to a person, not only the word code, but also the letter codes and feature codes are activated. This result is assumed to hold independently of

the particular word recognition model chosen (cf., Figure 4). But t is immediately presents a new problem. If all the familiar codes at all levels are activated, how do we select a particular code or unit for association to another cognitive system? How do we insure, when we are reading, that the highest level of perceived units is the one processed for meaning? When the first grade teacher holds up a card with a word on it, how can one know whether a child is attending to the letter units, or the spelling units, or the whole word?

Selecting Unit Levels

The issue of choosing a level of processing has been a major concern in our laboratory over the past several years. We have tested the processing of letters and familiar bigram units with a task which simply requires the person to match these patterns. For example, the subject is shown two letters positioned side by side, with several spaces between them, and is asked to press a button if the two letters match, but to withhold his response if they do not match. Sometimes we have him press another button when the two patterns do not match, but the data of interest are the latencies of correct matching operations. Sometimes we use novel letters as well as familiar letters, and sometimes novel letter-strings as well as familiar spelling patterns and words. The basic assumption of this method of matching is that we can measure the visual perception of a pattern relatively directly, as opposed to the more indirect method of measuring time to name a pattern or categorize it, since these latter tasks have other components which take up processing time, namely the association between the visual code and the code in the other system as well as the time to process the code in the other system (cf., Figure 1).

Consider a stimulus display containing two bigrams (cl cl) which the subject is instructed to observe. If the two patterns match, he is to press a button. According to Figure 1, the familiar letters c and l will be automatically activated and the familiar cluster unit cl will also be automatically activated. Notice that the subject could perform the task by matching each letter one at a time, involving two matching operations, or he could match the clusters as units which involve one matching operation. Which does he use? Even if he were sounding the letters and bigrams, something would have to determine which way is chosen. As it turns out, we have run several control conditions which indicate that the subjects' reaction times seem to be based on visual matching, even though they often report that they are aware of naming or pronouncing the patterns (cf., also Baron & Thurston, 1973; Pollatsek, et al., 1975).

Rohn Petersen and I (Petersen & LaBerge, 1975) set out to determine whether we could control the level of processing of bigram clusters. To do this we used two kinds of lists of items. One list contained predominantly familiar bigrams or clusters, such as sl, ph, sh, and br. The other list contained predominantly unfamiliar clusters such as ls, hp, hs, and rb. Our hypothesis was that for the first list, subjects would process the bigrams as units and make one matching operation, whereas for the second list, they would process the bigrams letter-by-letter and make two matches. In this way we hoped we could exert control over the unit level at which the stimulus display would be processed for a match.

What we needed was an indicator of the processing level used in each list. We inserted into these two kinds of lists two kinds of test items: a familiar bigram pair such as (cl cl) and an unfamiliar bigram pair such as (lc lc). We reasoned that, in the list of unfamiliar bigrams these two

test items should be processed similarly, that is, letter-by-letter. Therefore their latencies should not differ. However, in the list of familiar bigrams there should be a difference in the time to process these test items. The unfamiliar test item should take longer than the familiar test item because the subject is expecting to match the cluster as a whole, but cannot do it quickly because the cluster is unfamiliar. Even if he tries to match it letter-by-letter it should probably take time to switch attention to the letter level to make the match this way.

The results of this experiment are shown in Figure 2. The mean latency to

---

Insert Figure 2 about here

---

test patterns is plotted over two blocks of trials for each type of list. Familiar bigrams are labeled Clusters and unfamiliar bigrams are labeled Letters. When Cluster and Letter bigrams are embedded in lists of letters, there was no significant differences in the time to match them. But when the same Cluster and Letter bigrams were embedded in lists of clusters, there was a significant and substantial difference between them. Clearly the type of list context had an effect on the way a given stimulus pattern was processed. Thus we can take as our indicator of level of processing of the list items the difference in latency between unfamiliar and familiar test items. When there is no latency difference between the types of test items then we infer that the subject is processing the list of items from units at the letter level; when the unfamiliar items show a longer latency than the familiar items, then we infer that the subject is processing the list of items at a higher level.

The fact that the latency of matching an unfamiliar bigram changed with the type of list context tells us that list context must have been exerting some kind of influence on the processing of the units. Along with Estes (1974) we represent the influence of contextual conditions on perceptual processing by special contextual nodes, as shown in Figure 3. The arrows indicate

---

Insert Figure 3 about here

---

that a contextual node can activate the entire set of pattern codes at one level of processing. For example, when the letter contextual node is activated, it activates all the letter codes. This internal activation combines multiplicatively (Estes, 1974) with incoming sensory information so that individual letters appearing on the sensory surface are given emphasis compared to units representing groups of letters. This results in a higher activity of readout from the letter codes to the matching operation, so that the first match completed is made on the basis of letter units. We are not ruling out the possibility that readout from unit codes representing clusters would also reach the matching operation; but if they did they would reach it sometime after the letter code readouts arrived. In this way a context node selects the level of processing by which a given task operation is carried out. Typically selection can move upwards or downwards in the hierarchy of processing levels, although sometimes special training is needed to induce a child to read out from lower levels, as in segmenting words or syllables into phonemes. Again, it is important to note that Figure 3 represents a modification of the original LaBerge and Samuels (1974) model because, as it stood before, there was no mechanism by which a particular unit level could be selected when a string of letters is displayed to the subject. The change that is suggested here is

the addition of context nodes which combine with stimulus inputs to determine which level of units has highest priority for further cognitive processing.

If we tentatively accept a mechanism like the context nodes, we are immediately confronted with another problem. How is a particular context node itself activated? Something makes us look at letters when we proof-read a manuscript for spelling errors, and controls our processing of words and phrases when we proof-read a manuscript for sense, for example, to determine whether a sentence has been left out. Similarly, something controls the child's choice of units when the teacher holds up a card with the word "pen" and instructs the child to tell what the word means, in which case he processes it at the word level, or instructs him to tell how many letters it contains, in which case he will select a letter unit. In all of these examples, instructions determined the unit level selected. In the experiment just described, which called for matching units at different levels, the unit level was presumably determined by the list context. Exactly how experience with the first several items of a list produced an activation of a particular context code is not spelled out here.

We believe that questions concerned with the selection of context for a given task require research involving converging operations in order to provide clear and firm answers. As a step in this direction, we have carried out a series of experiments (LaBerge, Petersen, & Norden, 1976) to determine whether a subject can be instructed or cued to select a visual unit by direct methods. Using the same bigram materials in the experiment just described, we mixed the familiar and unfamiliar digrams within the same block of trials, but always presented familiar bigrams in the lower part of the screen and the unfamiliar bigrams in the upper part of the screen. Before each trial a circle was shown in the upper or lower position to tell the subject the

location and therefore the type of bigram to expect. We hoped that this very obvious spatial cue would induce the subject to process the items in the upper position letter-by-letter and the items in the lower position bigram-by-bigram. Using the same probe tests as before, we found no significant differences in mean latencies between the two presentation positions. This result apparently tells us that subjects could not easily switch levels of visual processing from trial to trial. It also suggests that it may take several trials to establish a stable context level. While much more research is needed to clarify this issue, it does appear that even college subjects have difficulty in directly shifting attention effectively and quickly to levels of processing of visual units. It appears that we cannot easily focus attention on a context in the way we can focus on a given pattern, or even on an operation such as adding or subtracting. It seems more likely that a subject can more directly select phonological levels because one can respond by pronouncing a letter or a syllable. Selection of a level in the phonological system might then feed back to the visual system and select the corresponding level of units.

Thus although there remains work to be done to solidify and extend these results, there are two conclusions from this discussion which seem to me to be of considerable import in the understanding of the role of units in perceptual processing. Firstly, since the subject has an option as to what unit he will process when he is shown a word, some source of information other than the immediate stimulus display must determine the unit selected. The second main point is that, at least for the visual system, the selection of unit level by direct deployment of attention to that level is not easily done by the average person.

Learning New Units

The foregoing discussion of the function of context nodes in the selection of perceptual units suggests strongly that the context node may also have a crucial role in the acquisition of new units. If an output from the perceptual system to a response is presumed to require the joint action of the information from the stimulus display and the activation by a context node, then we cannot expect to teach a child to respond on the basis of a new unit unless he already has its context node available and functioning. For example, if he is sounding a word on the basis of letter units, how can we expect him to shift to higher orthographic units unless there is control by the contextual node at this higher level? Thus, it seems quite likely that the formation of a context node precedes the formation of a unit at that level of processing.

It would appear then that we must address ourselves to two main questions concerning the learning or development of perceptual units: Firstly, how is a unit formed when a contextual node is in active existence? And, secondly, how is a unit formed when a new contextual node is not available and must be developed at the same time? On the one hand we learn new units most frequently at a level of processing at which we have already acquired some units, but on the other hand, it follows that we must occasionally learn a unit which is the first one acquired at a new level of processing. In this latter case, the child must acquire or develop the appropriate context node before or at the same time the first unit at that level is formed.

A very important issue in the perceptual learning of a word involves the regularities or rules of combinations of spelling patterns (Spoehr & Smith, 1973). For example consonant cluster units seldom follow each other; rather a vowel or vowel cluster usually intervenes. The present discussion does not

elaborate the relationships of visual units and orthographic rules, but rather concentrates on the implications of contextual nodes for the processing of simplified displays of various units. If this approach is successful in these situations, then hopefully a subsequent paper will give more discussion to orthographic rules.

Let us consider first how perceptual unitizing may occur at a level in which other units have already been formed. For example, how are letter units formed by children who have already learned a few letters of the alphabet? Or consider the unitizing of a new word by a child who already has several words in his visual vocabulary. For these children, we assume there exist context nodes for letter units and word units. Thus, all that is needed is to form a new unit of the kind with which he is already somewhat familiar. This is the type of situation in which we have studied the perceptual learning of novel letters (LaBerge, 1973) and others have studied the perceptual learning of new letter strings (e.g., Barron & Pittenger, 1974). This kind of perceptual learning has been conceived as having three stages (LaBerge, 1976), namely feature discovery, unit formation, and automatic unitization. For example, in the perceptual learning of artificial letters, the subject must first discover the distinctive features, much in the way E. J. Gibson's notions of distinctive-feature pick-up takes place in perceptual learning (Gibson, 1969). However, since the typical subject has already encountered many letters, he will find that selection of a single feature will not suffice to distinguish the new letter from the others he has learned. For example, suppose the new letter contains two lines which cross. Taken alone, this feature will lead to a confusion of the lower case letters t and f if the crossing lines are horizontal and vertical. Seldom does one feature distinguish one letter from all the rest of the letters which the person is called upon to distinguish.

This fact forces the learner to observe more than one feature in order to discriminate patterns from each other. When patterns are presented side-by-side in a same-different task, the subject can scan each contrasting feature independently, especially when the two patterns are positioned close to each other to facilitate scanning back and forth. But if he is presented one pattern at a time for identification, then the subject is required to learn the combination information of the two features. When we learn to identify letters by name we have to note the particular combinations of features by which each letter is uniquely identifiable. This stage of perceptual learning, therefore, requires that information in the distinctive features be combined. We consider this to be a second stage of perceptual learning in which a letter unit or code is formed.

We may consider the perceptual learning of a word in a parallel manner. For the four words at, it, an, and in, the task of merely discriminating pairs of these words can be carried out by scanning the first and second letter positions. For example, if a subject is given the pair (at an) to distinguish, he looks at the first letter position and finds a match, then he scans the second letter position and finds a mismatch, and then he responds indicating that the two patterns are different. But if he is required to identify (e.g., by four different responses) each of these words in isolation, he may scan each letter position, but now he must remember the first letter when he notices the second letter. It is the combination of the two letters which determines the response he will make. The perception of this combination information we assume to be the second stage of perceptual learning.

Now exactly how the combination information is processed is currently a very active controversy in perception. Perhaps the most well-known form of the question asks what information in a string of letters produces word

recognition. The vast array of experiments on the word advantage effect centers on this issue, since it asks how a letter is more easily detected when it is embedded in a word than in a non-word (e.g., Reicher, 1969; Estes, 1975). For a word, the particular combination (including order) of letters has been learned, but for a non-word, the combination information has not been learned, especially if it is a non-pronounceable non-word. Thus we consider this problem as falling under Stage II perceptual learning. It is clearly not a Stage I problem, since subjects have no difficulty discriminating a novel string of letters from other strings of letters.

The relationship of a word unit to its constituent letters perhaps can be clarified in terms of the network models shown in Figure 4. It is hoped

---

Insert Figure 4 about here

---

that this kind of classification of coding models will help to set priorities in research to determine what it is in graphemic stimuli that activates a word unit.

The Pure Hierarchy model assumes that letter code outputs converge to activate a word code, that is, a word is recognized by outputs from its constituent letter units. Of course, a word code may be activated by spelling patterns as well as by individual letters. However, for convenience in exposition of the six models, we have omitted the cluster level of units which we had included in Figure 1. To represent the Quasi-hierarchy model of word recognition, we add inputs to the word code from word features, such as word length, contour and internal relations. This model assumes then, that a word is recognized on the basis of component letters together with features unique to the word pattern.

The Two-level models assume that a word code is activated from the feature level only, so that the processing of a letter unit is not a necessary prior condition to the recognition of a word. The Two-level models differ with respect to the source of the features. In the Two-level Model A, some of the inputs to a word code come from features that also activate letters. In Two-level Model B, the word code is activated only by its own distinctive features, such as length, contour, and internal relations. In Two-level Model C, the inputs to a word code come only from features which also activate letter codes.

The five models just described assume that a word unit is activated by outputs from either letter codes or feature detectors, and that usually at least two of these inputs are required to produce the word code. The One-level model, on the other hand, assumes that each unit extracts its information directly from the sensory surface, without intermediating codes. In this way there is no fusion of two or more inputs, but rather a direct relation between the stimulus information and the unit.

It will be noted that the One- and Two-level models do not have the hierarchical property whereby the codes of one level feed into the codes of the next level. The arrangement of units in these models are stratified in Figure 4 only to reflect the objective hierarchical ordering of features, letters, and words. I would like to point out again that the inclusion of the spelling-pattern level which is critical for the functioning of orthographic relationships is not expected to alter the main contrasts being drawn among these six models. If the spelling pattern units were added to Figure 4, the additional crossings of lines would produce a rather noisy array which would tend to obscure the main differences among the models.

Perhaps we can summarize the main principles of these models by borrowing on variations of an old Gestalt principle. For the Pure Hierarchy Model, the word is the sum of its letter parts. For the Quasi-hierarchy Model, the word is more than the sum of its letter parts. For the One- and Two-level models, the word is different than the sum of its letter parts.

Thus, there appear to be at least six different ways that a word code might be activated once the stimulus information falls on the sensory surface. We are not suggesting that only one model is true for all cases. It may be that we shift from one processing model to another under different reading demands; and it may well turn out that, in the course of learning to read, a child shifts from one strategy to another strategy as the predominant way of processing a word.

If we assume a Pure Hierarchy or Quasi-hierarchy Model, then one interpretation of unitizing is that a word is fused from component letter units or some feature units that are unique to the word. But if we assume the Two-level B Model, then it might be said that the word unit is fused from its own unique set of features. In this case Stage II perceptual learning is not based on previously formed units, but rather on a reselection of features, such as particular internal relations in a word. This process may be considered quite close to the process of Stage I perceptual learning, since it involves discovery of the appropriate features. However, it differs in at least three respects from the search for distinctive features. Firstly, the purpose here is not usually that of discriminating two patterns from each other, since this can already be easily done on the basis of the known letter units. Secondly, the word features are likely to be relational features that are not point-to-able in the sense that letters can be singled out. This implies the third difference, namely that the reselection is not likely to

involve a search by re-orientation of the eyes to other locations within the word pattern, but must, in some sense, be guided to a new relational feature based on the raw information already existing in the print of the letters. After all, if the ink marks were removed the word could not be perceived. Therefore there is something about the marks on the page which gives rise to word perception as well as to letter perception.

Regardless of whether Stage II unit learning is regarded as a fusion of letters or features, or as a direct extraction process, or some other process, it would seem that there must be a shift in the context node from activating units at the letter level to activating units at the word level.

Once the subject perceives the word as a unit, further exposures are needed to consolidate the unit so that eventually the unit can be activated without attention to fusion and/or reselection of features. In this way perception becomes automatic. If Stage II unit learning for words proceeds by fusion of letter code outputs, then in Stage III the fusion occurs without the attentional scanning of the letters. If Stage II occurs by an extraction of relational features, then during Stage III the deployment of attention to this process becomes less and less.

Now we turn to the more difficult issue of learning the first unit at a new level of processing. I hope the reader will keep in mind that my comments in this section of the paper are speculative. However, it is in view of the practical importance of initial unit learning that I stray further from supporting data. Hopefully, the ensuing remarks will stimulate some productive tests of these notions.

The situation of learning a new type of unit is much like that confronting a child who has no trouble identifying or naming letters of the alphabet, but now must perceive a cluster of letters, e.g. the bigram "ch" or "sh" and

notice the regularities of their combinations which is involved in what some mean by learning orthographic "rules". These letter clusters must be perceived as units to associate effectively a unitary phonological response to them. Similarly after phonics training with parts of words, the child must learn to blend the parts into a word. This may involve shifting context levels in phonological system. Listening to a child trying to blend a word for the first time can be very instructive. For example, in blending the two sound patterns "ch" and "at" into /cæt/ the child's pause between /c/ and /æt/ indicates quite directly that he is processing these as separate units. Later, after "something clicks in his head", he pronounces the word "chat" in a quick, uninterrupted manner, which seems to the observer to indicate that he has processed the word at the phonological level holistically. We can be even more convinced of this unitization if he recognizes the meaning of the word at the time that he successfully seems to blend it. The point is that once the child has learned to blend successfully once or twice, the blending of new cases seems to proceed faster. Similarly, when the child learns to identify one or two bigrams as units, he seems to pick up new bigrams and even trigrams, etc., quite quickly. The problem appears to be located at the point of getting the child to identify a bigram as a unit or blend parts of a word for the first time. We often say that he has to "learn the skill" of blending or unitizing; this is perhaps not unlike the learning-to-learn phenomenon described by Harlow (1949). In terms of the present model, we say that when he can activate the appropriate context node to perform the appropriate unitizing, he has learned an "acquisition skill". By means of this skill he goes on to acquire new units at that level.

Now some might remark that the appropriate context nodes of letter groups and whole-word sounds are available and actively used by the child at the time

he first learns to blend and unitize letter strings. His daily conversation employs spoken words, and he can almost always recognize his own written name, and usually a few common written words. Doesn't this mean that he already has an active node for these particular perceptual units and therefore for all other units at these levels? I would venture to guess that the child does not necessarily have the appropriate context nodes for other units because the context is not the same. Consider how the child recognizes his written name among other names presented to him. Suppose John E. is written above a coat hanger in the Kindergarten classroom. John's task is to discriminate the visual pattern of the written name from other names he sees. In a typical group of 20 or 30 children in his class, there is not likely to be another written name very similar to his. If there is a John F. or a Joan E. he might show some confusion, and a teacher often foresees this problem and places the labels at different spatial locations so that the child can use spatial cues to help him discriminate the names. If pressed, I would suggest that the context node involved here is at the feature level. Thus generally, the young child can identify his own written name on the basis of a few features that are quite different from the information he must learn to pick up when he identifies the vast range of words in his visual vocabularly two or three years later.

The parallel argument for phonological units is more difficult and admittedly somewhat tentative at this time. But I would ask you to consider that before a child learns to read, he does not segment spoken words very often. A sentence is heard as a continuous stream of sounds, and boundaries are often not at all distinctive to a child until after he has learned to read. Therefore, the way he hears and utters words in normal speech probably involves altogether different contextual controls than those appropriate to

the sounding of words in isolation. Although one has seen and heard his own name many times before he begins to read, he may eventually come to hear and see the name differently after he has learned to read. When I say "see and hear differently", I mean that a different context node is involved, and this implies that he must be able to pick up stimulus information from the same words differently.

Let us return to the case in which the child has learned to identify individual letters and now is called upon by task requirements to process strings of letters as units, e.g., bigrams or short words. His training with letters has sensitized him to the locations within a word which carry the information which distinguishes one word from another. For example, there are many words in our language which differ by only one letter. A child who has focused attention on letter units must now free himself from this narrow range of information and take in new information whose source is no longer individual letters, but rather is in the combination information in the string of letters. As we have pointed out before, there are several viewpoints as to how the information in a string of letters gives rise to a new unit representing a spelling pattern or a word. Regardless of the view adopted, the word unit is assumed to be controlled by an appropriate contextual node. And this node must either be present genetically through maturational readiness or somehow acquired by appropriate experiences. If it is available, then the child learns the new unit and moves rapidly ahead. If it is not available, we either wait until it matures or we try to induce the context by instructional means.

Now, although scientists prefer to control the events in their domains as opposed to waiting for them to happen, and although our practical educational objectives also press us to control beneficially the learning of a

student, it may turn out that the development of a contextual mechanism is not subject to direct control and therefore may be resistant to efforts to "teach" it. Recall the experiments just described in which we tried to cue the level at which a subject would match pairs of bigrams. The results indicated that subjects apparently do not directly activate levels of processing, implying that the context nodes may be so remote from attentional control that efforts to teach a subject to use them, when they have never before used them, may be entirely fruitless. It seems to me that the notion of readiness captures the flavor of this problem. If we have to wait for a child to become ready to learn a skill, this implies that the basic processing mechanisms are not receptive to our attempts to control them. This does not mean that some skills already in the child's repertoire could not substitute for the one in question. A child who does not perceive the bigram "sl" as a visual unit may stay at the visual letter level and combine them phonologically by sounding letters one after another. On a more general level, a reader may compensate for an inability to perceive many whole word units by sounding out syllables.

If control of visual unit levels is not easily accomplished directly, then perhaps other routes to the activation and selection of visual units are possible. In the experimental example described, in which the familiarity of a list of bigrams apparently determined the level of processing, there must have been some pathway by which the visual nodes were selectively activated. Our best guess is that the operation of matching the patterns produced a feedback to that level of units which yielded the faster latency of matching. The ultimate control of level of units in this case was the environment, since it was the type of display which determined the context. Similarly, we often find that a child can "get into the swing" of performing a new skill under the strong control of classroom stimulation, but when faced with the same task

alone, cannot perform the new skill. What seems desirable is to provide the child with a means of controlling his own level of perception rather than leaving it to the preceding sequence of stimulation.

One promising direction for providing internal control of visual processing levels looks to the phonological system, in which control is typically exerted by speech. Adults find it relatively easy to listen either to whole words or the last rhyming syllable, or the sound of the first phoneme. This may be done quickly because the person can sound these units to himself. In this way he selects a level of processing. Now if he has already learned sound-to-sight associations, i.e., he can categorize sounds as words, clusters, or letters, then he can control the level of visual processing by attentional activation of appropriate sound units. According to this line of reasoning, phonics skills may be critical for control of visual processing levels. Children who have not learned to segment phonemes within a word and learn the relations between sound levels and sight levels should therefore have considerable difficulty in selecting the appropriate visual units in the graphemic displays presented to him. Thus, the acquisition of new visual units, especially at a new level, is regarded here not as a simple selection and/or fusion of sensory inputs from the bottom-up. Because of the important interactive role of the contextual nodes with incoming stimulation, perception and acquisition of new units requires that appropriate activation of contextual nodes be controlled. The major direction of this control is assumed to come from higher cognitive systems. Therefore the acquisition of a new unit is regarded here as based both on sensory information and contextual mechanisms which are largely under the control of higher cognitive systems.

It may be the case that fast learners differ from slow learners because of a difference in the availability of contextual mechanisms. We have already

seen evidence that slow and fast learners apparently do not differ with respect to speed of forming associations (Zeaman & House, 1963; Estes, 1970), but rather with respect to perceptual learning. Specifically, lower mental age subjects require more time to discover the appropriate perceptual feature or dimensions of the perceptual patterns in a task. Similarly, one might expect to find differences in learning rates at other stages of perceptual learning, such as coding new units. Perhaps one of the most dramatic differences between individuals may be in terms of the speed at which they grasp new unit levels. Some children move from letters to higher orthographic units to whole phrases so rapidly that it is difficult sometimes to believe that for these children there are separable subskills at all. But for others, the contextual levels may become available relatively slowly and we are then painfully aware of a particular level because for a time the child cannot move from that level to the next.

I have been talking about the learning of units in the same way that some researchers refer to the learning of skills. In view of the similarity of our conceptualizations of skill learning and unit learning we should not be surprised if we find that they share common questions. For example, it is reasonable to assume that a global skill such as reading can be appropriately segmented into subskills for the prupose of effective instruction. If so, does the most effective way to teach these skills proceed from the bottom-up, that is, beginning with smaller units and proceeding to the larger units? If we find ourselves spending a great amount of time on some particular subskill, how do we motivate the child appropriately? For example, if the motivation for reading is to comprehend written language, how do we most effectively bring this motivation down to the subskill learning at the level of orthographic units?

Let us ask to what extent unitizing skills are teachable. I believe that instruction, whether by a teacher or by a book, exercises control over the perception of the student most effectively by presenting material that is point-to-able. The teacher can make sure that the student picks up the critical information when he says "consider this" or "look at that", whether it is a feature of an airplane, a note of music, a sentence, or a mathematical expression. When a point-to-able item is the critical feature, then learning apparently proceeds quickly, because the "pointing" function of the teacher controls the attentional focusing of the student. But, when the feature is not so easily pointed out, then attentional focusing of the student is not directly controlled by the teacher, and there will be uncertainty about the material the teacher intends him to perceive. Many relational features are of this type. For example, when two pitches are presented simultaneously, it is not clear that the student perceives the two individual notes or the particular musical interval. Similarly when the teacher points to the word "man", it is not clear that the student perceives the visual relations within the word. Compare the words "man" and "mat". One is tempted to say that the difference is perceived in terms of the different terminal letter unit because we can point to the terminal letters n and t. But it is more likely that the reader is perceiving the whole word, and the first two letters, "ma", and the last letter are interrelated. How does a teacher effectively point to such a relation? Indeed, it is far more convenient to point to the letter than to a relation. And this may mean that it is very difficult to "point the way" to new levels of perceiving, especially when they involve relational features such as this example represents.

However, once a level has been brought under the control of a contextual mechanism, then the teacher may have several means of promoting the acquisition

of units. Two of those methods have already been described. The first way uses units which are already at the new level to induce the context by presenting them as a series of examples before presenting the new unit. For example, if the bigram "dr" is to be learned as a unit, it can be presented following "sl", "fl", and "gr", which we shall assume are already familiar units at that level. The second way of promoting unitizing would use sound-to-sight relations to induce the appropriate visual level. For example, sounding the series of visually presented words, "mink", "sink", and "pink" and then "drink" may promote perceiving "dr" as one visual unit, if the child already can segment the initial phoneme and the initial letter in the first three words.

If, however, the child has not visually unitized a bigram prior to "dr" nor has learned to visually segment the first letter or letters of a word and associate sounds to these visual units, then teaching the child to unitize "dr" visually will be more difficult. If the teacher points to "dr", the child is likely to perceive two letters. If the teacher pronounces it, "Jruh", the child may not isolate the initial phoneme from the neutral syllable "wh", but, more importantly, he may be hearing the two phonemes "d" and "r" separately. If he hears two sounds, then he may look for two symbols, and, although this may be desirable in some tasks, it is not desirable here for the purpose of promoting visual unitization.

While the foregoing examples are quite rough attempts to illustrate ways that unitizing might be taught, it is hoped that they do provide a contrast between teaching a new unit when the appropriate contextual control of a level is available and when it is not available.

It is obvious that considerable research is needed to fill in the gaps in the contextual control modifications of the LaBerge-Samuels model (1974).

It is difficult to formulate optimal algorithms for acquiring units when there is considerable uncertainty as to how a word code is activated. Figure 4 shows six different models, and there are undoubtedly others. Quite probably the acquisition of word codes proceeds by stages which may be arranged hierarchically. Initially, the child may recognize a word in terms of letters and spelling patterns, and only later as a single unit. But if the learning sequence does proceed hierarchically, this does not necessarily imply that when the word is finally recognized as a unit that it is processed hierarchically, i.e., from outputs from letter and spelling pattern codes. Fluent readers, in fact, probably learn new words directly as single units without going through letter and spelling pattern stages, owing to use of contextual nodes of a high order. Before we can begin to probe these questions we need to have reliable indicators of the level at which a word is processed in any given instance. Figure 2 illustrates one way this might be done, but the test has drawbacks in terms of ease of administration.

Another difficulty standing in the way of prescribing ideal conditions for acquiring units is the lack of detailed knowledge of the way contextual information interacts with sensory input. Furthermore, as mentioned before, it is not clear how visual contextual nodes are activated by the phonological system and perhaps the syntactic-semantic systems.

If we were to compare the importance of learning a unit at a new level with learning a unit at a familiar level at which other units of the same type have already been acquired, there would be no question that the accomplishment of the first case is the more momentous educational event. A contextual jump presents a new class of units within the grasp of the child; and we often refer to these moments as times when the child has "made a leap", or is "over the hump". When he has learned to use a new contextual node, he

has in effect learned an "acquisition skill", by which a host of new units at that level can be acquired. It may turn out that contextual jumps are events which we cannot directly point out for the child, but rather must happen to him, and all we can do is provide the best conditions under which these happenings can occur. Defining what those conditions are may be a fruitful direction of research in our efforts to understand the role of perception in the very complex skill we call reading.

In summary, the process of reading begins with the perception of visual patterns. Whether the unit of the pattern selected for deeper processing is a letter, a spelling pattern, a word, or a word-group depends upon an interaction between the stimulus display and the dominant contextual mechanism. These mechanisms are presumably driven by "top-down" connections from operations on patterns (e.g. matching), phonological events, and semantic-syntactic events. Details of the influence of higher-order cognitive systems upon output selection from the visual system remain to be worked out. Meanwhile, there may be value in speculating how the development of a given contextual mechanism could affect the speed of acquisition of a new unit.

References

Baron, J., & Thurston, I. An analysis of the word-superiority effect. Cognitive Psychology, 1973, 4, 204-229.

Barron, R. W., & Pittenger, J. B. The effect of orthographic structure and lexical meaning on "same-different" judgments. Quarterly Journal of Experimental Psychology, 1974, 26, 566-581.

Broadbent, D. E. Perception and communication. London: Pergamon Press, 1958.

Estes, W. K. Learning theory and mental development. New York: Academic Press, 1970.

Estes, W. K. Memory, perception, and decision in letter identification. In R. L. Solso, (Ed.) Information Processing and Cognition: The Loyola Symposium. New York: Wiley, 1974.

Estes, W. K. The locus of inferential and perceptual processes in letter identification. Journal of Experimental Psychology: General, 1975, 104, 122-145.

Gibson, E. J. Principles of perceptual learning and development. New York: Appleton-Century-Crofts, 1969.

Gibson, E. J., & Levin, H. The psychology of reading. Cambridge, Mass: MIT Press, 1975.

Harlow, H. F. The formation of learning sets. Psychological Review, 1949, 56, 51-65.

LaBerge, D. Attention and the measurement of perceptual learning. Memory and Cognition, 1973, 1, 268-276.

LaBerge, D., & Samuels, S. J. Toward a theory of automatic information processing in reading. Cognitive Psychology, 1974, 6, 293-323.

LaBerge, D. Perceptual Learning and Attention: In W. K. Estes (Ed.) Handbook of Learning and Cognitive Processes, Vol. 4, Hillsdale, NJ: Lawrence Erlbaum Associates, 1976, in press.

LaBerge, D., Petersen, R. J., & Norden, M. Exploring the limits of cueing. In S. Dornic (Ed.) Attention and Performance; Vol. VI, London: Academic Press, 1976 (in press).

Mandler, G. Organization and memory. In K. W. Spence & J. T. Spence (Eds.) Psychology of Learning and Motivation, I. New York: Academic Press, 1967.

Melton, A. W., & Martin, E. (Eds.) Coding processes in human memory. New York: Wiley, 1972.

Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 1956, 63, 81-97.

Moray, N. Attention: Selective processes in vision and learning. London: Hutchinson, 1969.

Petersen, R. J., & LaBerge, D. Contextual Control of Letter Perception. Technical Report #15, Minnesota Reading Research Project, University of Minnesota, Minneapolis, Minnesota, 1975.

Pollatsek, A., Well, A. D., & Schindler, R. M. Familiarity affects visual processing of words. Journal of Experimental Psychology: Human Perception and Performance, 1975, 1, 328-338.

Shiffrin, R. M., & Gardner, G. T. Visual processing capacity and attentional control. Journal of Experimental Psychology, 1972, 93, 72-82.

Spochr, K. T., & Smith, E. E. The role of syllables in perceptual processing. Cognitive Psychology, 1973, 5, 71-89.

Treisman, A. Selective attention in man. British Medical Bulletin, 1964, 20, 12-16.

Zeaman, D., & House, B. J. The role of attention in retardate discrimination learning. In Handbook of Mental Deficiency. New York: McGraw-Hill, 1963.

## Footnote

This research was supported by United States Public Health Service Grants HD-06730 and in part by the Center for Research in Human Learning through National Science Foundation Grant BMS-75-03816. The author acknowledges the helpful discussions with S. Jay Samuels, Rohn J. Petersen, and Lee Brownston.

## Figure Captions

Figure 1. Representation of units in the visual, phonological, and semantic-syntactic systems. Solid dots represent familiar memory codes which can be activated automatically from the sensory surface. Open dots represent unfamiliar codes which require additional activation by attention. Lines connecting codes represent inter-system associations.

Figure 2. Mean latency and percent errors of matching responses for familiar digram (Clusters) and unfamiliar digram (Letters) test items for two types of list contexts.

Figure 3. Representation of units in the visual system for three levels of processing. Contextual nodes activate all units at the indicated processing level.

Figure 4. Six coding models for word recognition. Solid dots represent familiar memory codes at three levels of processing (Spelling pattern codes are omitted here -- convenience of illustration). Context nodes activate all units at the indicated level of processing.

VISUAL SYSTEM
SENSORY SURFACE
feature
letter
cluster
word
m (word)
SEMANTIC – SYNTACTIC SYSTEM
PHONOLOGICAL SYSTEM
p (letter)
p (cluster)
p (word)

38

39

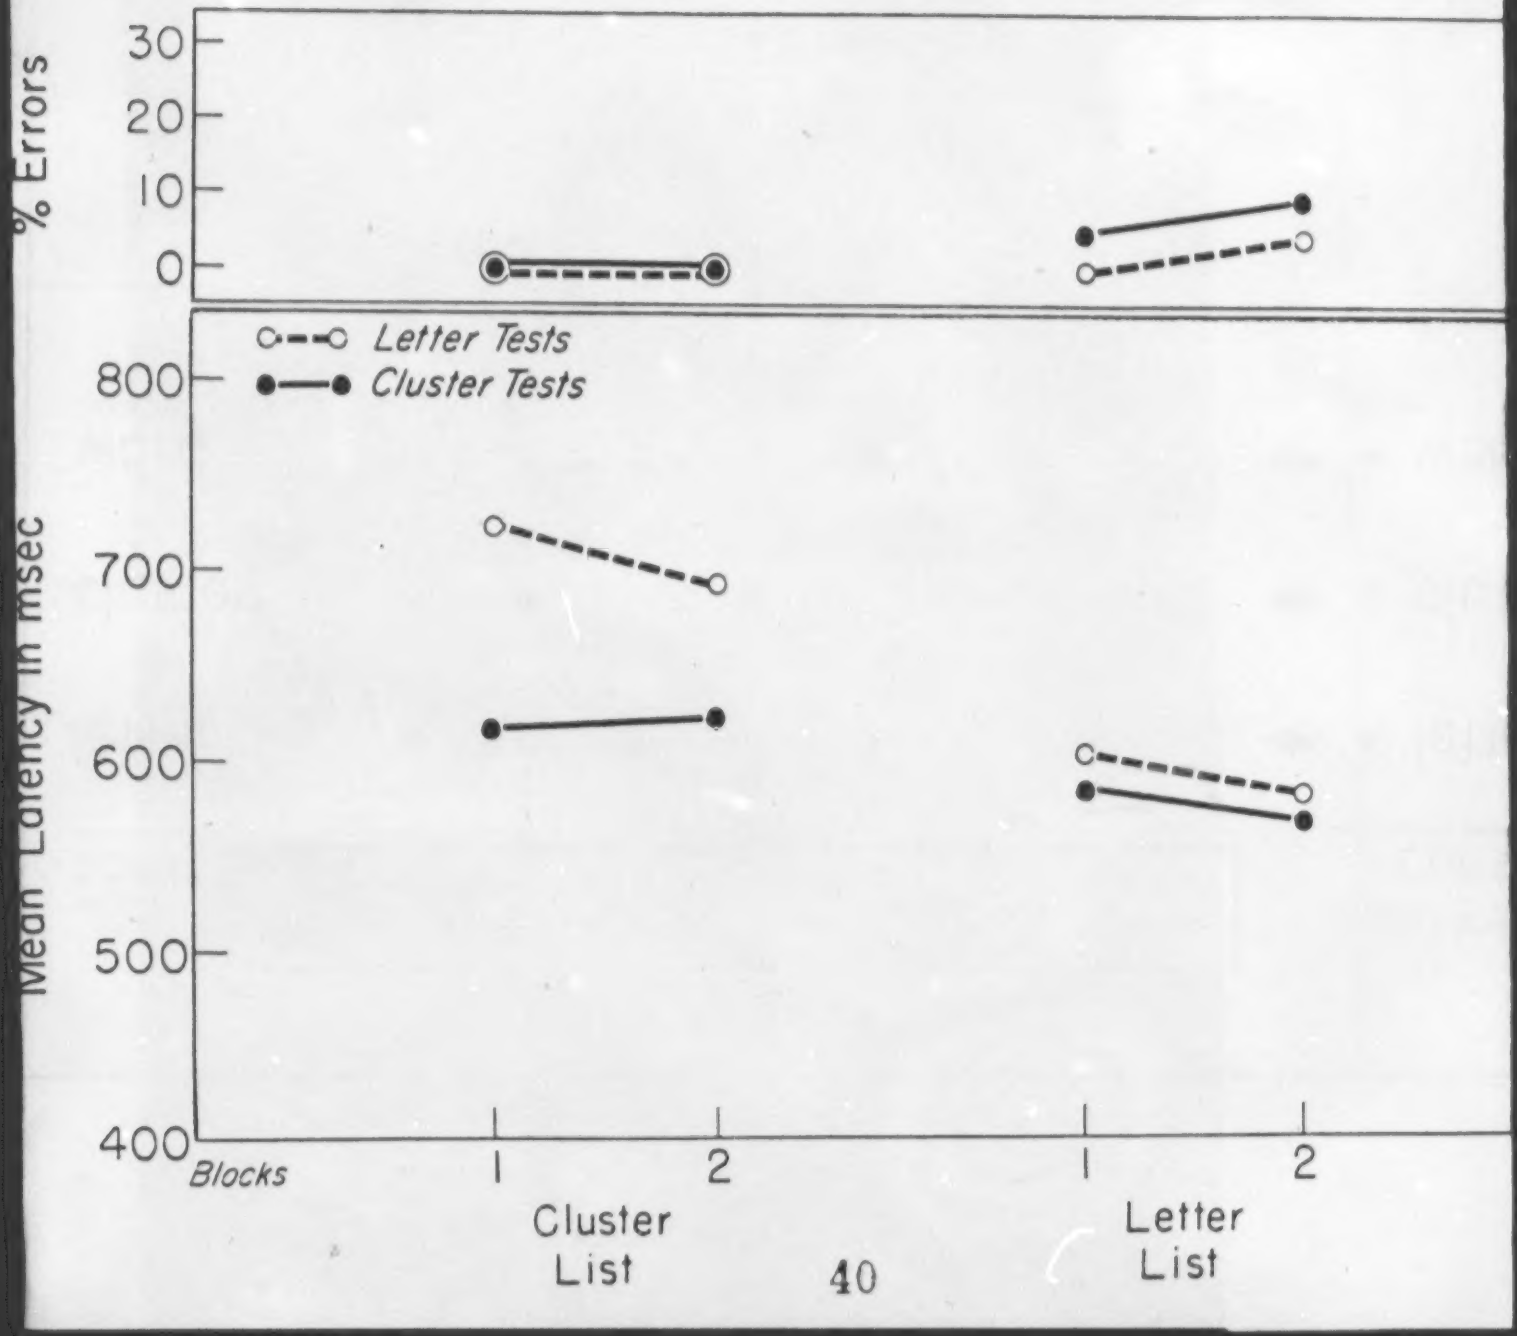
% Errors
30
20
10
0
Letter Tests
Cluster Tests
Mean Latency in msec
800
700
600
500
400
Blocks
1
2
1
2
Cluster List
Letter List

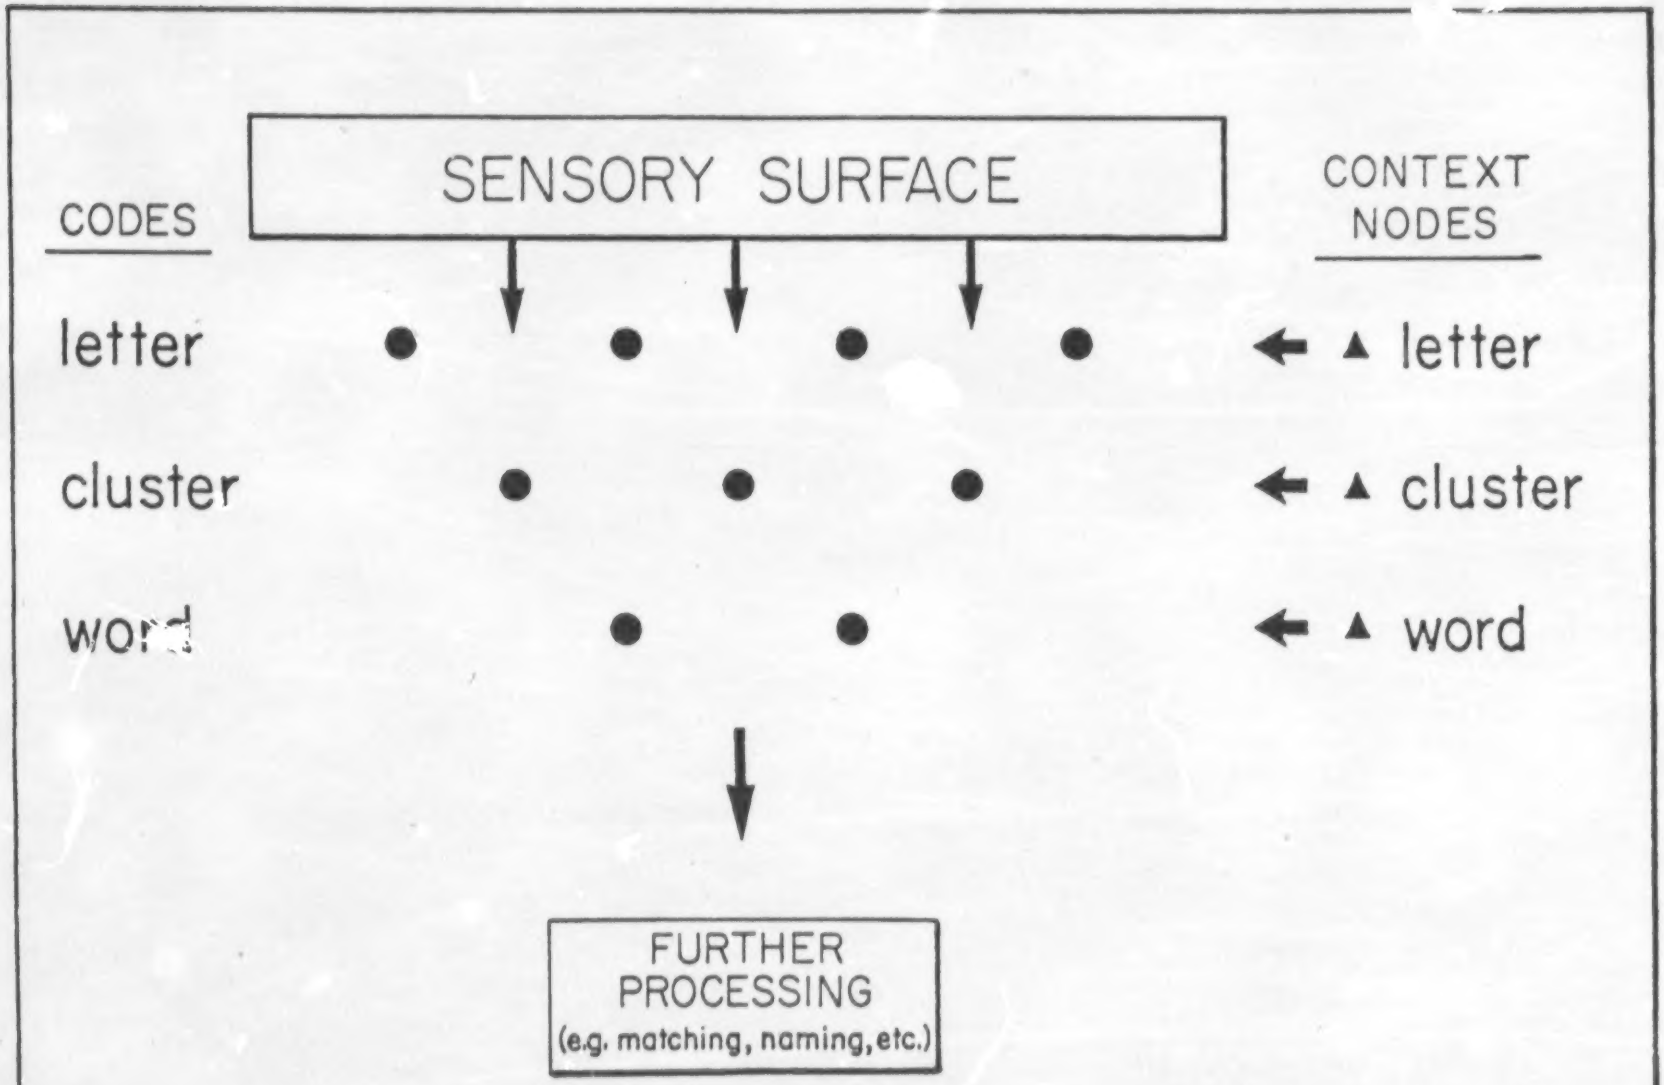
SENSORY SURFACE
CODES
CONTEXT NODES
letter
cluster
word
letter
cluster
word
FURTHER PROCESSING
(e.g. matching, naming, etc.)

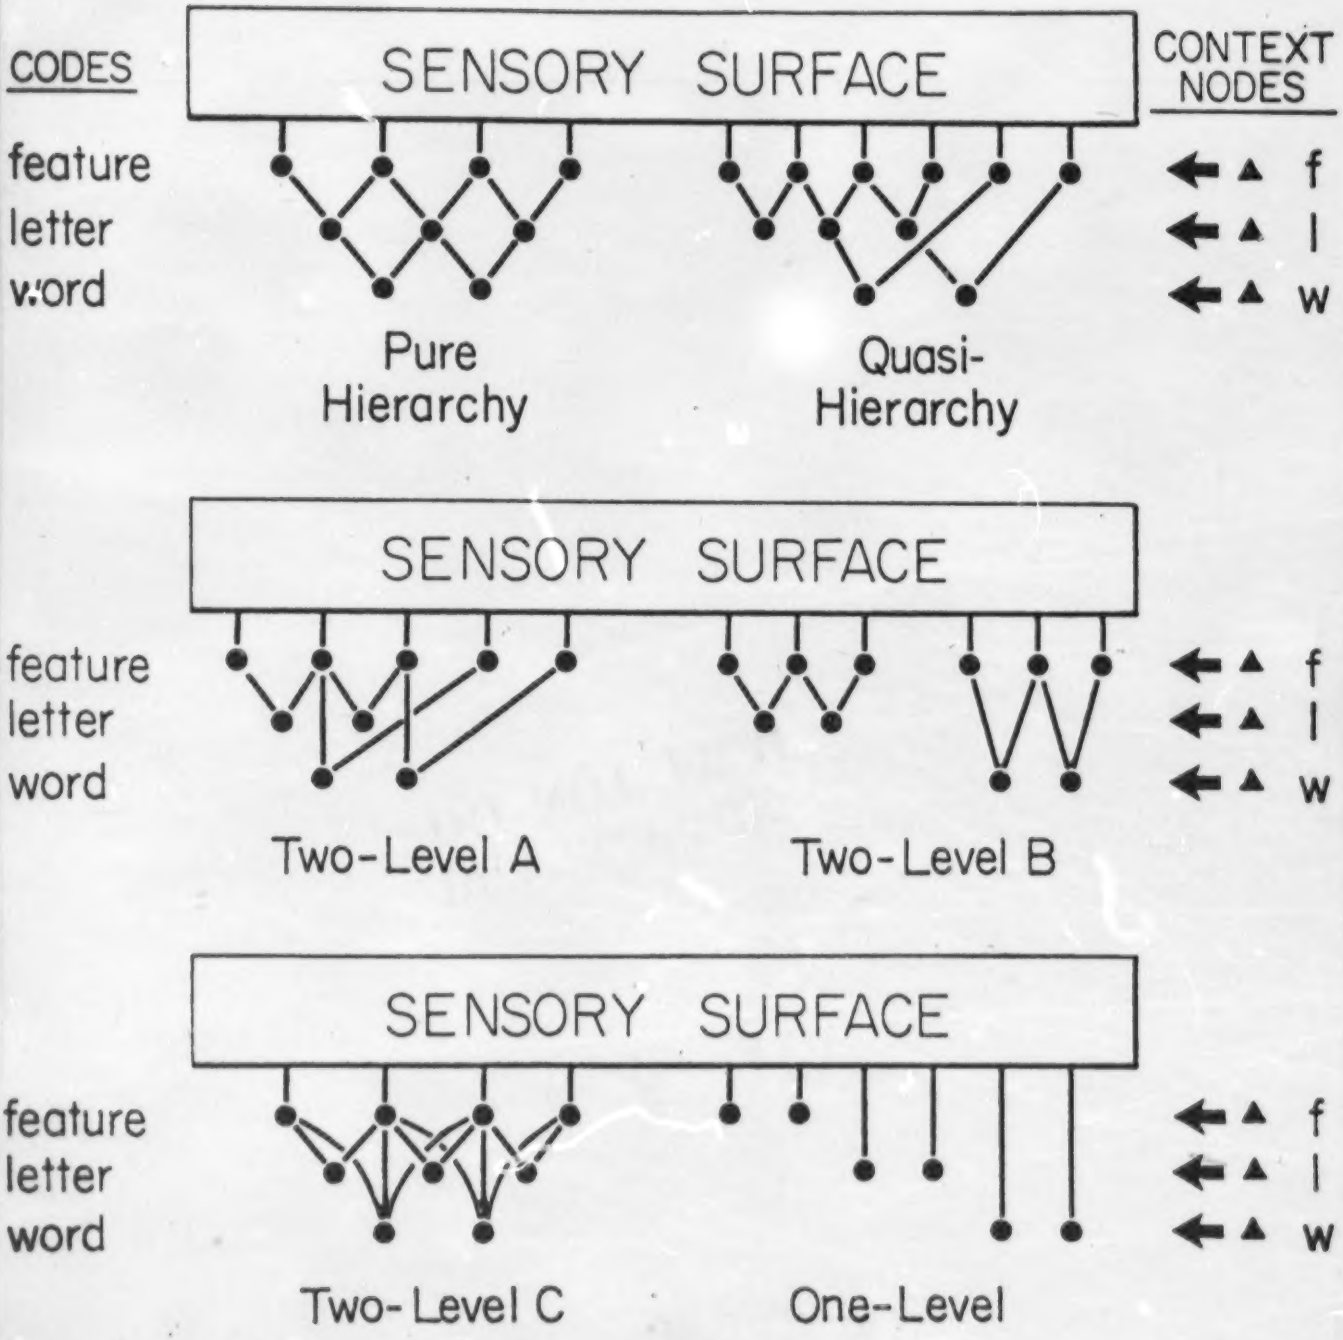
CODES
SENSORY SURFACE
CONTEXT NODES
feature
letter
word
f
l
w
Pure Hierarchy
Quasi-Hierarchy
SENSORY SURFACE
feature
letter
word
f
l
w
Two-Level A
Two-Level B
SENSORY SURFACE
feature
letter
word
f
l
w
Two-Level C
One-Level

## OPEN DISCUSSION OF LaBERGE PRESENTATION

GREGG: Why does perceptual learning take longer than associative learning? Why do you want so many stages of perceptual learning, and all of those questions?

LaBERGE: Associative learning is all-or-none and fast. Perceptual learning involves discovery of features and unitizing, which is relatively much slower. If you try to learn four names to those four things (indicating), you will find that the problem is not so much in assigning a name to four objects, it is in identifying uniquely those four objects. The difficult problem in the letter sequence BDPQ is not association of B to that pattern and D to that pattern and P to that pattern and Q to that pattern. The problem is getting yourself straight on the relations among the features, and that takes a long time.

GREGG: Are you saying that once familiar objects can be recognized, then you can use them to form the next level of association? It seems to me that your six models of word recognition are at different levels.

LaBERGE: Possibly, yes.

GREGG: I think that you got it backwards, Dave. You don't decide what the unit is first, you learn what the patterns are from the redundancies of the language, from the experience, from the extent to which the teacher is able to control attention to larger and larger parts of the word.

LaBERGE: Yes. I certainly am not saying that the only way to learn is from the bottom up. I am just saying that if you are putting anything through the system,

these models show six ways that it can go through the system. The features of a pattern have to be activated before the person can see anything further. I am suggesting this maps out alternatives.

Perhaps I am not really providing a specific answer to your question; what I am doing is trying to clarify what I consider a very muggy situation in our discussions; namely, as was brought out before: "Let's set out what the alternative models are, and let's do some experiments to decide among them."

A student of mine has a dissertation that has tested these six models, and surprisingly comes out very clearly and strongly for one of those.

WALLACH: No doubt there are readiness processes, but don't you think that readiness processes should be a hypothesis of last resort, adopted only if one has really tried everything one can think of instructionally, and has to admit failure?

LaBERGE: Yes. I should like to make sure that I didn't give the impression that we had final evidence to show that we have, once and for all, determined what we can do in controlling the child's learning, with respect to these levels of processing. I don't want to say that at all. I just want to say I have experiments which indicate that for college students it is very difficult to get at the levels of processing, and probably something that looks like readiness is taking place. And being an experimental psychologist, I want to have as much under my control as I can, and as teachers you want to have as much as you can, so I am certainly on the same side as you are.

WALLACH: I just want to point out there is a background issue that lurks here,

that has to do with the social context. And readiness can be a code word for "Wait until the child can do it, and lo and behold middle class children can do it sooner." And that has implications.

LaBERGE: Yes. Of course, in the wrong ears, and the unexperienced ears, that statement could be used counterproductively. That is certainly very clear. I don't want to say, "Leave the kid alone, and he will read."

RESNICK: Another way of bridging what looks like an argument is to recognize that even though the child has to do the induction and make the leap, there are different ways of organizing stimuli, organizations that make that leap more likely, and other ways to make it less likely.

LaBERGE: Exactly.

RESNICK: I think that's what the instructional level has to be about, and that was explored only on the surface.

LaBERGE: Consider an experiment by Art Roneti, Samuels' student, which I think is related to this question. He overtrained children with letter recognition and letter naming, and found that it interfered with the children's moving up to the next level. They just simply stuck at that level too long. If you train your children to be very conscientious, and above all avoid any errors, they will never risk taking a leap to a larger unit. And if they don't ever take that risk, they simply will never go above some particular level, at least not for a long, long time. This is one condition which would work against encouraging a child to go up to a higher level. I don't think there is any difference in our

opinions here; but I think there may be a difference in the way you feel I should state things. Is that it?

WALLACH: I think there may be a difference in how much one tries to do different kinds of things.

LaBERGE: I think it is important to determine this with some rather clearly articulated models. I am not saying these models are really clear and beautifully testable. But I believe that this is at least a step in that direction, and that you get some clarification of what you are talking about. Then you can answer questions like this one, and then perhaps find the kinds of research methodology which will determine a real breakthrough, where we could actually control something like this.

TRABASSO: I have the impression that you are focusing by and large on a visual processing model.

LaBERGE: Yes.

TRABASSO: I am wondering, since we are talking about early reading, to what extent you want to bring in auditory codes, and the role of the auditory system? It seems to me that there are some correspondences in the auditory codes which relate to the levels. For example, words are large units, and are difficult to decompose into smaller units. One might be able to decompose them into, say, syllables and decompose them into phonemes. And so there is some correspondence in the auditory system to what you are representing in the visual system. Okay?

LaBERGE: Yes.

TRABASSO: The question is when does unitization takes place. It might be that the auditory code operates and is not necessarily reducible to lower-level codes, or single-letter, visual codes. Okay?

LaBERGE: Sure.

TRABASSO: My question is: If you are going to extend this model to early reading, how would you bring in the auditory system?

LaBERGE: I would be tempted to do the following: Because you can pronounce the phonological units, you have more attentional control over what level you work with. I can pronounce letters and I can pronounce these clusters. But I could draw these things visually as clusters or draw them as letters, but that isn't as firm a control over the level as my speaking them.

That's why I emphasize the visual level at this point, separated from the auditory level, because I think there is a difference in terms of the control you can have over the level of processing. By the kinds of response I make, I can put myself at a level pretty quickly. Then given you are at that level, you could have the feedback from the phonological system to these three, four, five, or six visual contextual levels. Remember, I said attention apparently can't get at them directly. But there must be some indirect way to do it.

Then presumably if the child can name the word as he looks at it, it may be that that kind of naming would help him to move up a level in the visual system. I don't have facts, but I would be encouraged to try the following: When I

wanted to teach a child to begin to look at something at a higher level, let's say orthographic units, I would not train him on things which are close to it, which you do when you are trying to do feature discovery at an early perceptual level. I would try to get him to name it, hoping that this would feed back to the visual levels to control the level of processing. So in identification or naming you get him to say the whole thing, using something that he knows already; or if you are on the phrase level, try to get him to take in these big units. Find some familiar phrase he can speak, and then use the identification response to train the level in the visual system.

That is speculation. We know there is a lot of research that needs to be done. It is complicated; you have to go through the forest in some sort of orderly way, and I am suggesting this is one way to try to order it. I can see an experimental design coming out of that, can't you?

GLASER: I hope the experimental design you see coming up is one in which you try to produce the best and most serious possible training conditions, to follow the study where you didn't get results by training. Seriously, there are a lot of training conditions that you haven't tried.

LaBERGE: That's right. I will say what I have tried it on myself as a subject. In the "CL, CL, CL, CL," and "LC, LC, LC and LC," experiment, I speak the letters to myself, but data do not come out the way you would expect. But we haven't done that extensively, for I may be a bad subject.

LESGOLD: Why are you taking the experiment you have described as really strong evidence for a lack of ability to switch levels. It seems that in the experiments that you mentioned, it may even be optimal to switch levels,

depending on how long it takes to switch attention from one level of analysis to another, one might be better off ignoring the information. An experiment that might come closer is one in which there was, say, two- or three-second warning, with an arrow pointing at the top.

LaBERGE: That's exactly what we have, two seconds' warning.

LESGOLD: So the subject knows then for sure?

LaBERGE: Yes.

RESNICK: Would your comments still hold?

LaBERGE: Yes, my comments hold but we haven't tested the limits of it. We have tested within the same kinds of boundaries governing the cueing a person to get ready for a particular pattern or a particular operation. And within those limits cueing levels do not work. So that's why I say this is more difficult than the ordinary kind of cueing.

POPP: It is interesting that all of your evidence comes from college students, and that perhaps the responses that appear to have been built up in the experimental situation may already exist. To take your results and generalize to how you would teach a beginning child, who is not already familiar with dealing with the CL clusters and the LC, may be risky. I wonder if one way of relating to the beginner would be to look for differences in specific clusters. For instance, one might hypothesize that the same-different discriminations of your CL versus LC may be greater than the DR versus RD, because the RD exists in the

language of a final consonant cluster and LC is a rare final cluster.

LaBERGE: I was careful about that when we put it up, we didn't use RD.

POPP: Well, there may be others. For instance, LS also occurs as a final cluster. My point is that the college students already have been exposed to clusters over a period of time. To then say we can't train the small child to do it, seems to me to be a bit of a leap.

LaBERGE: I see what you mean.

POPP: Did you say we can't train the child to switch, or to move up a level?

LaBERGE: I said that it is out of attentional control. If it is not directly controllable by the college student's attention, then I could hardly hope that it would be controllable in a direct way by the child. We have run this standard experiment on children and got a significant difference with early first graders.

POPP: Early first graders?

LaBERGE: Early first graders, apparently before they formed units. But I think your point is well taken. There is a leap there, but I tried to take you across that leap with the use, hopefully, of clearly articulated theoretical notions.

DANKS: You seemed to imply that there is a single node for each letter, which is associated with pronunciations or the name. We might need a separate node for the visual properties of a letter versus its name because the name is not really

used in reading for the most part. I think it is better represented in Lee Gregg's model, where he had an alpha for a single node, and an A associated with it. Maybe we need two separate nodes, one with the visual properties and reading pronunciations, but not letter names, and then a second node with the name of the letter, the alphabetic sort of thing.

LaBERGE: Good point. In fact, it has been our intention to separate letter names from letter sounds, so you have two places it can go. That in itself requires a context node to determine which will operate in a particular case.

SUPPES: Dave, could you say a little bit more about the evidence that is in favor of the two-level model? Because of the way you described it, it isn't clear how you would discriminate between them. Indicate how you would discriminate between the two of them. I realize that is a hard question.

LaBERGE: I could sketch it very, very quickly to you. This is Rohn Petersen's dissertation. What he did was to take advantage of the pathway activation notions of Michael Posner and Snyder, in which it is assumed that if you had presented a word to a person, and the hierarchy model was true, then you would be activating the letters. Because anything that activates the word has to come through the letters. So these letters would be hot.

But in the two-level model, when you activate the word, you are not necessarily activating the letters. So what he did was the following: He showed the person a word like "golf," and if it was a word, he had to get ready to see the word "four," if it was a nonword, he had to get ready to see "five." Okay?

Occasionally, following a word such as "golf," we would present to him anagrams for matching, like "flog, flog," and occasionally we would give him things which had no letters in common, like "read, read." We said that if the letters here are activated when he sees the word "golf," (which he must decide is a word on the basis of the task we give him) then the anagram "flog" is going to be faster than the control word "read." If, on the other hand, the letters are not being activated, when he sees the word "golf," then the anagram "flog" will be seen as quickly as "read." What we found, again and again, is that the latency to "flog, flog" is not significantly different from the latency to the control, "read, read." It is clearly different, indicating that the letters here are evidently not being activated.

RESNICK: Can you show under some other condition that there is a difference between "flog, flog" and "read, read?"

LaBERGE: We have been trying to do that. We tested the subject occasionally on letters, like two L's. What we did though was have a person scan the array looking for a letter. If a letter L is there, press the left button, if the letter D is in there, press the right button. Now he is looking for each one of the letters. What we get is an attenuation of this difference, but it still occurs.

We also give him nonwords. If a person sees four graphemic things together, like that, whether they are words or nonwords, evidently the college students, at least, chunk it, and if he chunks it, then he tends to act according to the two-level model, so that he is not coming up through the letters.

So what these results indicate to us consistently, is that word recognition does not follow a hierarchical model in which letters must first be processed. It seems that words have their own emergent features which form the basis of recognition by fluent readers.

CARROLL: How do you know there are features coming through there? I understand your notion of what is happening at the word level, but how do you extrapolate that to get features?

LaBERGE: How do we get to the features? We are saying here there are only word features coming through. Because remember, here, the features are shared. Even though the features don't go through the letter level here, they still are shared with these levels, and if these features of the letters were shared, it should help him. So the features are part of the whole, as well as the letters are part of the whole, in the most rigorous statement of that model.

FARNHAM-DIGGORY: What would you predict for a beginning reader?

LaBERGE: Our discussion section indicates that something like this may be going on: The beginning reader processes a word letter by letter, and indeed our experiments indicate that the latency of recognition is a function of the length of the word. The reader is processing letter by letter, and when there is a slope. When there is no slope you know he is chunking. And of course, even if you give adults letter strings, not words, you will find a clear slope. We have clear evidence for that.

GREGG: I just objected to your statement he will have chunked it. He didn't

chunk it, it was chunked, he recognized the whole word.

LaBERGE: I am sorry, yes.

GREGG: You see what I mean?

LaBERGE: There is something I don't understand about it.

GREGG: There is something about the way you express things that have the arrow of time going in the wrong direction.